Robust variability of grid cell properties within individual grid modules enhances encoding of local space

William T. Redman[1,2,†], Santiago Acosta-Mendoza[1,†], Xue-Xin Wei[3,4,5,6], and Michael J. Goard[7,8,9]

[1]Interdepartmental Graduate Program in Dynamical Neuroscience, University of California, Santa Barbara
[2]Intelligent Systems Center, Johns Hopkins University Applied Physics Lab
[3]Department of Neuroscience, The University of Texas at Austin
[4]Department of Psychology, The University of Texas at Austin
[5] Center for Perceptual Systems, The University of Texas at Austin
[6]Center for Theoretical and Computational Neuroscience, The University of Texas at Austin
[7]Department of Psychological and Brain Sciences, University of California, Santa Barbara
[8]Department of Molecular, Cellular, and Developmental Biology, University of California, Santa Barbara
[9]Neuroscience Research Institute, University of California Santa Barbara
[†]Equal contribution

Abstract

Although grid cells have become one of the most well studied functional classes of neurons in the mammalian brain, the assumption that there is a single grid orientation and spacing per module has remained untested. We investigate and analyze a recent large-scale recording of medial entorhinal cortex to characterize the presence and degree of heterogeneity of grid properties, within individual modules. We find evidence for small, but robust, variability and hypothesize that this property of the grid code could enhance the ability of encoding local spatial information. Performing analysis on synthetic populations of grid cells, where we have complete control over the amount heterogeneity in grid properties, we demonstrate that variability, of a similar magnitude to the analyzed data, leads to significantly decreased decoding error, even when restricted to activity from a single module. Our results highlight how the heterogeneity of the neural response properties may benefit coding and opens new directions for theoretical and experimental analysis of grid cells.

The discovery of grid cells in medial entorhinal cortex (MEC) [1] has led to considerable experimental and computational work aimed at identifying their origin [2–12] and their function [13–21]. The organization of grid cells into discrete modules [1, 22, 23], with particular grid properties (grid spacing and orientation) conserved within module, but not between modules, has fundamentally shaped this research. For instance, the increasing size and spacing of grid modules along the dorsal-ventral axis of MEC, by discontinuous jumps of a near constant ratio, has been argued to be optimal for encoding local spatial information [24, 25], when grid cell activity across all modules is integrated together [24–28]. This is despite the fact that the hippocampus, a downstream target of the MEC, receives inputs from only a portion of the dorsal-ventral axis [29]. The modularity of the grid system has also been proposed to simplify the wiring necessary for generating continuous attractor dynamics [2, 3, 5, 16], a computational mechanism theorized to underlie grid cells function that enjoys considerable experimental support [23, 30–34].

Much of this prior theoretical work has made the additional assumption that grid cell properties are identical within a single module. However, the distributions of measured orientation and spacing show non-zero variability [22, 30, 34]. This could be due to finite recording time, neurophysiological noise, or the sensitivity of numerical methods used to fit grid properties. Alternatively, this variability could reflect underlying inhomogeneity, at a fine scale, within modules. Despite the fundamental way in which grid cell function has been informed by their modular organization, to the best of our knowledge, no characterization of the degree and robustness of variability in grid properties within individual modules has been performed.

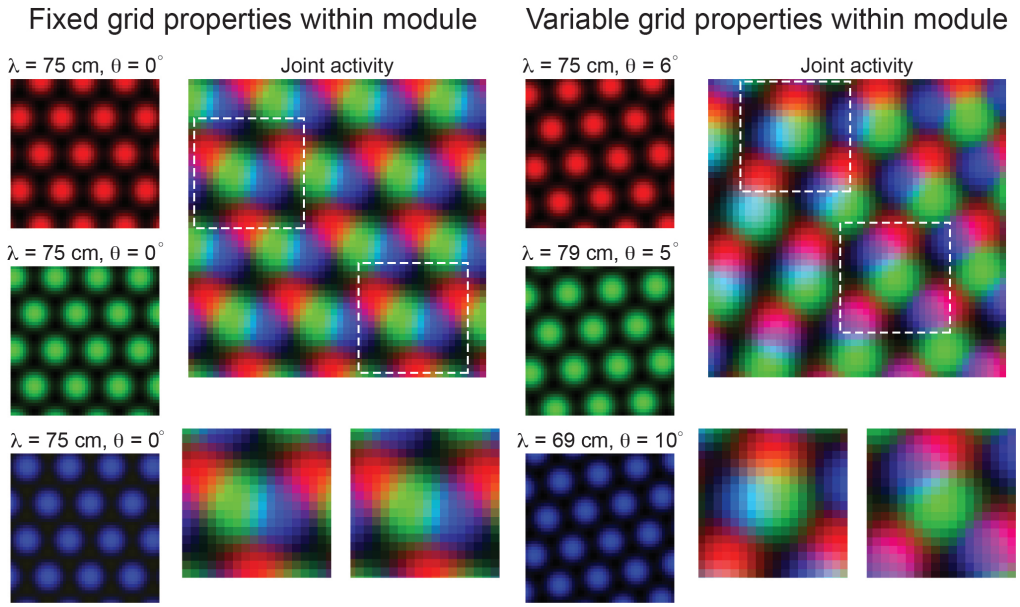

Figure 1: **Variability in grid cell properties, within a module, leads to joint coding of local space.** When the activity of three idealized grid cells, all with the same grid spacing and orientation, are considered, the periodicity of the responses limits the amount of information conveyed about local space (Left column – "Fixed grid properties within module"). That is, there are multiple locations in physical space with identical population level activity. However, when three grid cells with variable grid spacing and orientation (in the realm of what is measured within individual grid modules – see Results), their joint activity contains considerably more information (Right column – "Variable grid properties within module"). This spatial inhomogeneity would be expected to increase further with larger populations of grid cells. Dashed squares in the joint activity map are enlarged below.

If robust variability of grid properties does exist, then it is possible that this heterogeneity could be exploited to encode additional information about local space, reducing the requirement of integration across multiple grid modules (Fig. 1). In particular, when all grid cells have the same grid orientation and spacing, the joint population activity has a translational invariance that persists, even as larger populations of grid cells are considered (Fig. 1, "Fixed grid properties within module"). In contrast, if grid cells within a module have variability in their grid orientation and spacing, then – over a *finite* area – the translational invariance of the population activity is broken, and distinct patterns emerge in distinct locations of space (Fig. 1, "Variable grid properties within module"). As an example, the red and blue grid cells in Fig. 1 overlap most strongly in the lower portion of the arena (denoted by purple pixels), while the blue and green grid cells overlap most strongly in the left half of the arena (denoted by cyan pixels). This is despite the fact that the variability in spacing and orientation is only on the order of a few centimeters and degrees, respectively.

In this paper, we perform detailed analysis of recent state-of-the-art MEC electrophysiological recordings, which were made publicly available [34]. We characterize the variability of grid orientation and spacing within individual modules, and find evidence for small, but robust, variability. This variability is present whether a single value of grid orientation and spacing is assigned to each grid cell, or whether grid distortion [35, 36] is taken into account by considering three grid orientations and spacings independently. To assess the functional implications of this heterogeneity, we perform simulation experiments with synthetic, noisy grid cell populations, where we have complete control over the distribution of grid orientation and spacing. We find that the variability in grid cell orientation and spacing, at a similar degree as present in the data we analyzed, leads to lower decoding error of local space, when using the activity of a *single* module.

Taken together, our results challenge a long-held assumption and support a growing appreciation of the local spatial information encoded by grid cell populations [37–40]. Additionally, they encourage consideration of the broader benefits that multiple modules may provide, beyond the encoding of local space [41–43].

Results

Robust differences in grid cell properties within individual modules

To determine the extent to which variability of grid properties in individual modules exists, and to what extent this variability is a robust property of the grid code, we analyzed previously published data [34]. This data represents a major advance in MEC recordings, with tens to hundreds of grid cells simultaneously recorded, and allows us to characterize the distribution of grid properties, within a single grid module, to an extent not possible with other data sets.

For each grid cell, we compute the grid spacing (λ) and orientation (θ) by measuring properties associated with the six hexagonally distributed peaks of the spatial autocorrelogram, as traditionally performed [44] (Fig. 2A; see Materials). For clarity, we begin by focusing on a single module, recorded from an open field environment (recording identifier: Rat R, Day 1, Module 2 – R12). This module was picked for its long recording time (approximately 130 minutes of recorded activity, as compared to the other open field recordings that have 30 minutes or less of recorded activity); and for its large number of simultaneously recorded putative grid cells ($N = 168$). We find that in this module, high quality grid cells (grid score $>$ 0.85, $N = 74$) have a range of grid orientation and spacing (example cells shown in Fig. 2B; distributions across module shown in Fig. 2E, F), with λ ranging from 65 cm to 90 cm and θ ranging from 2° to 9°. Overlaying the spatial autocorrelogram of pairs of grid cells with similar grid spacing and different grid orientation (Fig. 2C), or vice versa (Fig. 2D) enables visualization of the extent of this variability.

To assess whether the heterogeneity of grid properties present in a single grid module is a robust feature or attributable to noise (either in the recording or the grid property fitting procedure), we measure the variability in grid orientation and spacing **within** a single grid cell and **between** pairs of grid cells. If the heterogeneity is explainable by noise, then we expect that the within-cell variability will be of the same size as the between-cell variability. In contrast, if the heterogeneity is a feature of the grid code, then we expect the within-cell variability will be significantly smaller than the between-cell variability.

To measure the within- and between-cell variability, we split the recording into evenly spaced 30 second bins, randomly assigning each temporal bin to one of two equal length halves, and recomputing the grid properties for each half of the recording (Fig. 3A; see Materials). We set inclusion criteria to filter out cells that did not have robust hexagonally structured spatial autocorrelograms across splits of the recording (see Materials). While these requirements were strict (see Fig. S1 for percent of cells rejected), they set a conservative estimate on the amount of grid property variability, ensuring that we did not artificially inflate the variability due to weak grid cells. We did not find that the length of the temporal bin used to split the data had a large impact on the percent of cells accepted by this criteria (see Materials, Fig. S2).

We find that the distribution of within-cell variability of grid orientation and spacing is more concentrated around 0 than the between-cell variability, across all 100 random shuffles of the data into two halves. Comparing the average within- and between-cell variability of grid spacing and orientation reveals that nearly all of the grid cells ($N = 82$) contained more between-cell than within-cell variability (Fig. 3D, E): 95.1% grid cells for orientation ($\overline{\Delta\theta}_{\text{within}} = 1.2°$, $\overline{\Delta\theta}_{\text{between}} = 2.0°$) and 100% of grid cells for spacing ($\overline{\Delta\lambda}_{\text{within}} = 1.9$ cm, $\overline{\Delta\lambda}_{\text{between}} = 6.9$ cm). A Wilcoxon-Signed-Rank Test indicated that between-cell variability was significantly higher than within-cell variability for both orientation and spacing ($\Delta\theta$: $p < 0.001$, $\Delta\lambda : p < 0.001$)

In keeping with convention, the reported grid cell properties are the average of those computed for each of the three independent axes in the spatial autocorrelogram (Axis 1: aligned to $\approx 0°$; Axis 2: aligned to $\approx 60°$; Axis 3: aligned to $\approx -60°$ [45]). To ensure that this averaging was not contributing to the greater between-cell variability, we repeated the analysis above, restricting ourselves to each axis separately. The results again demonstrate that grid properties are significantly more robust within-cell than between-cell (Fig. 4A, B). For every axis, the average between-cell variability for every cell was significantly higher than the average within-cell variability (Fig. 4C, D), as reported by the Wilcoxon-Signed-Rank Test for orientation ($p < 0.001$ for all three axes) and for spacing ($p < 0.001$ for all three axes)

Having demonstrated that grid cell properties are robustly heterogeneous in a single module, we proceeded to analyze the remaining recordings in the data set ($N = 420$ cells across 9 modules; one recording had

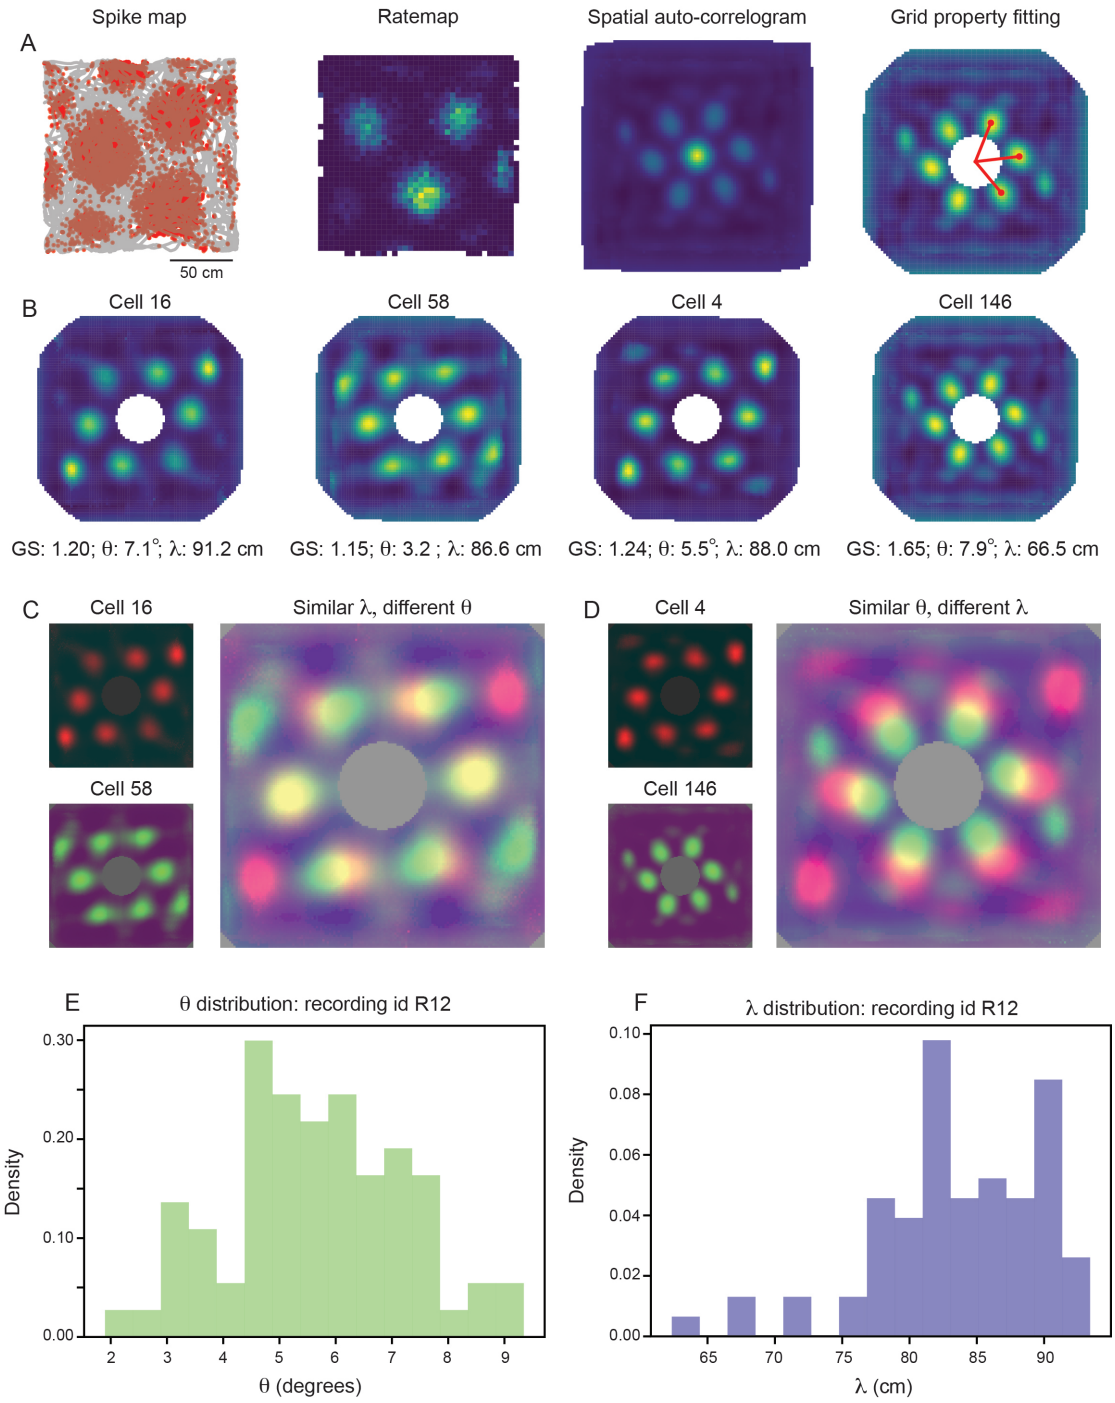

Figure 2: **Grid properties are variable within a single grid module (recording ID R12).** (A) Overview of the standard procedure used to calculate the grid spacing and orientation of a given grid cell. First, spike maps are computed by identifying the location of the animal at the time of each spike. Gray lines denote trajectory of rat, red dots denote location of spikes. A rate map is constructed by binning space and normalizing by the amount of time the rat spent in each spatial bin. A spatial autocorrelogram is computed, and after the center peak is masked out (white pixels in the center of the spatial autocorrelogram – leading to change in color scale), the grid properties are fit by measuring the length and angle of the three closest hexagonal peaks. (B) Example grid cells from the same module (recording ID: Rat R, Day 1, Module 2 – R12), with estimated grid score, orientation (θ), and spacing (λ). (C)–(D) Spatial autocorrelograms overlayed for two pairs of grid cells [from (B)]; one pair had different θ and similar λ (C), and the other has similar θ and different λ (D). (E)–(F) Distribution of θ (E) and λ (F) across all grid cells in recording, with grid score > 0.85.

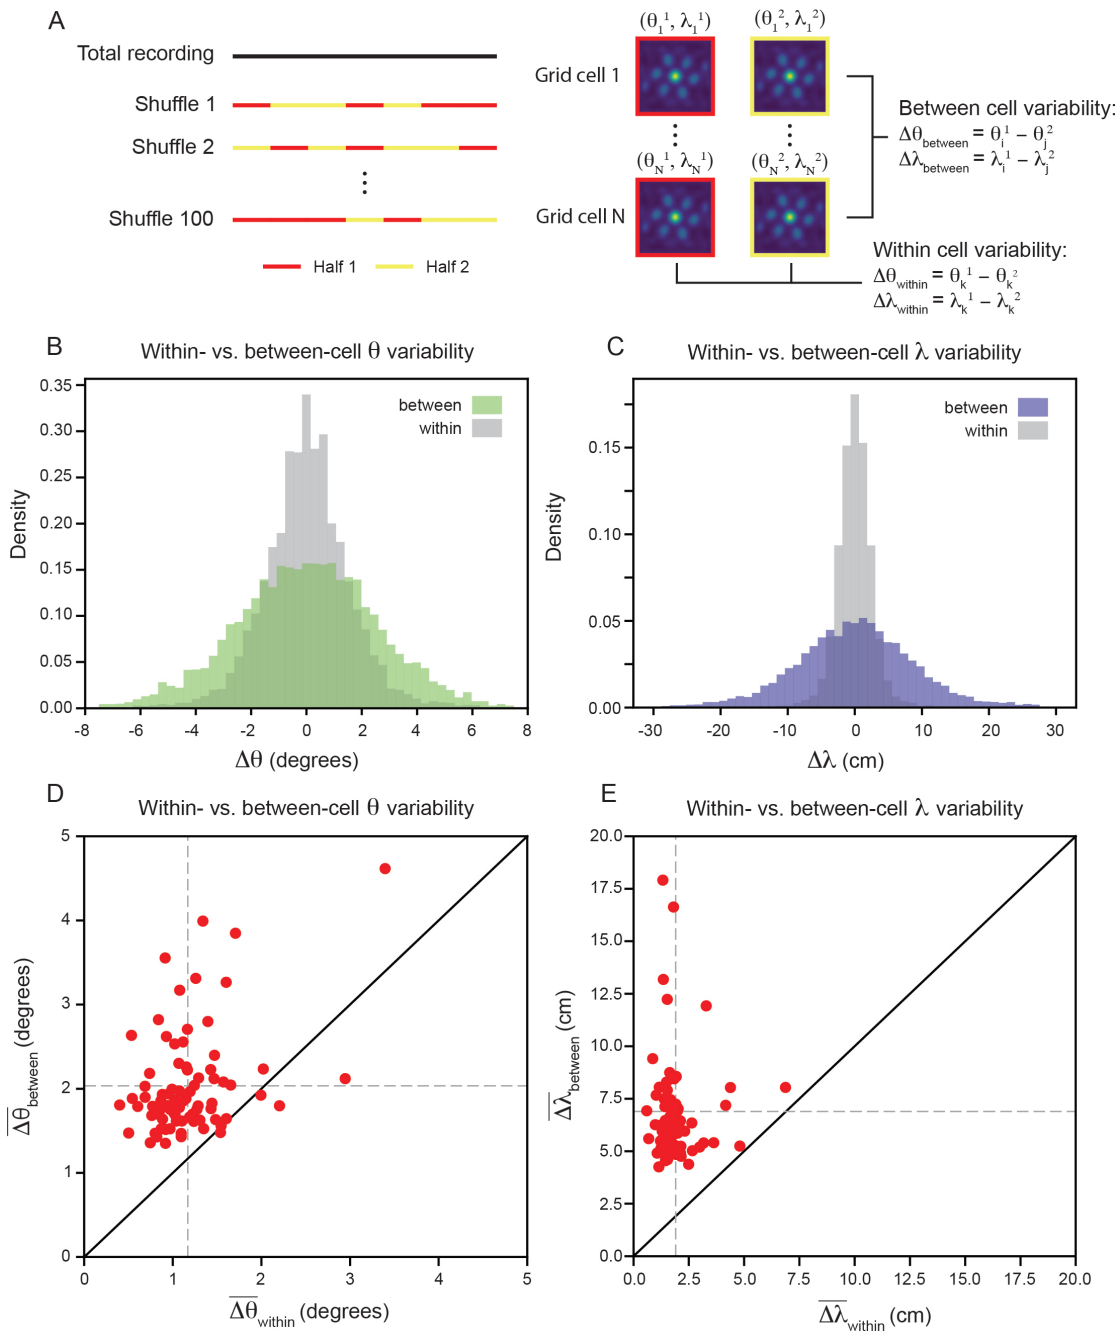

Figure 3: **Variability of grid properties is a robust property of individual grid module (recording ID R12).** (A) Schematic overview of approach used to compute the between- and within-cell variability of grid orientation and spacing. (B)–(C) Distribution of within- and between-cell variability of θ and λ, respectively. Note that the distribution is across all 100 random shuffles of the data into two halves. (D) Average within-cell variability of grid orientation ($\overline{\Delta\theta}_{\text{within}}$), compared to average between-cell variability of grid spacing ($\overline{\Delta\theta}_{\text{between}}$). (E) Same as (D), but for λ. 1 cell was excluded from (E) for visualization ($\overline{\Delta\lambda}_{\text{between}} > 20$ cm), but was included in non-parametric statistical analysis. $N = 82$ in (D)–(E).

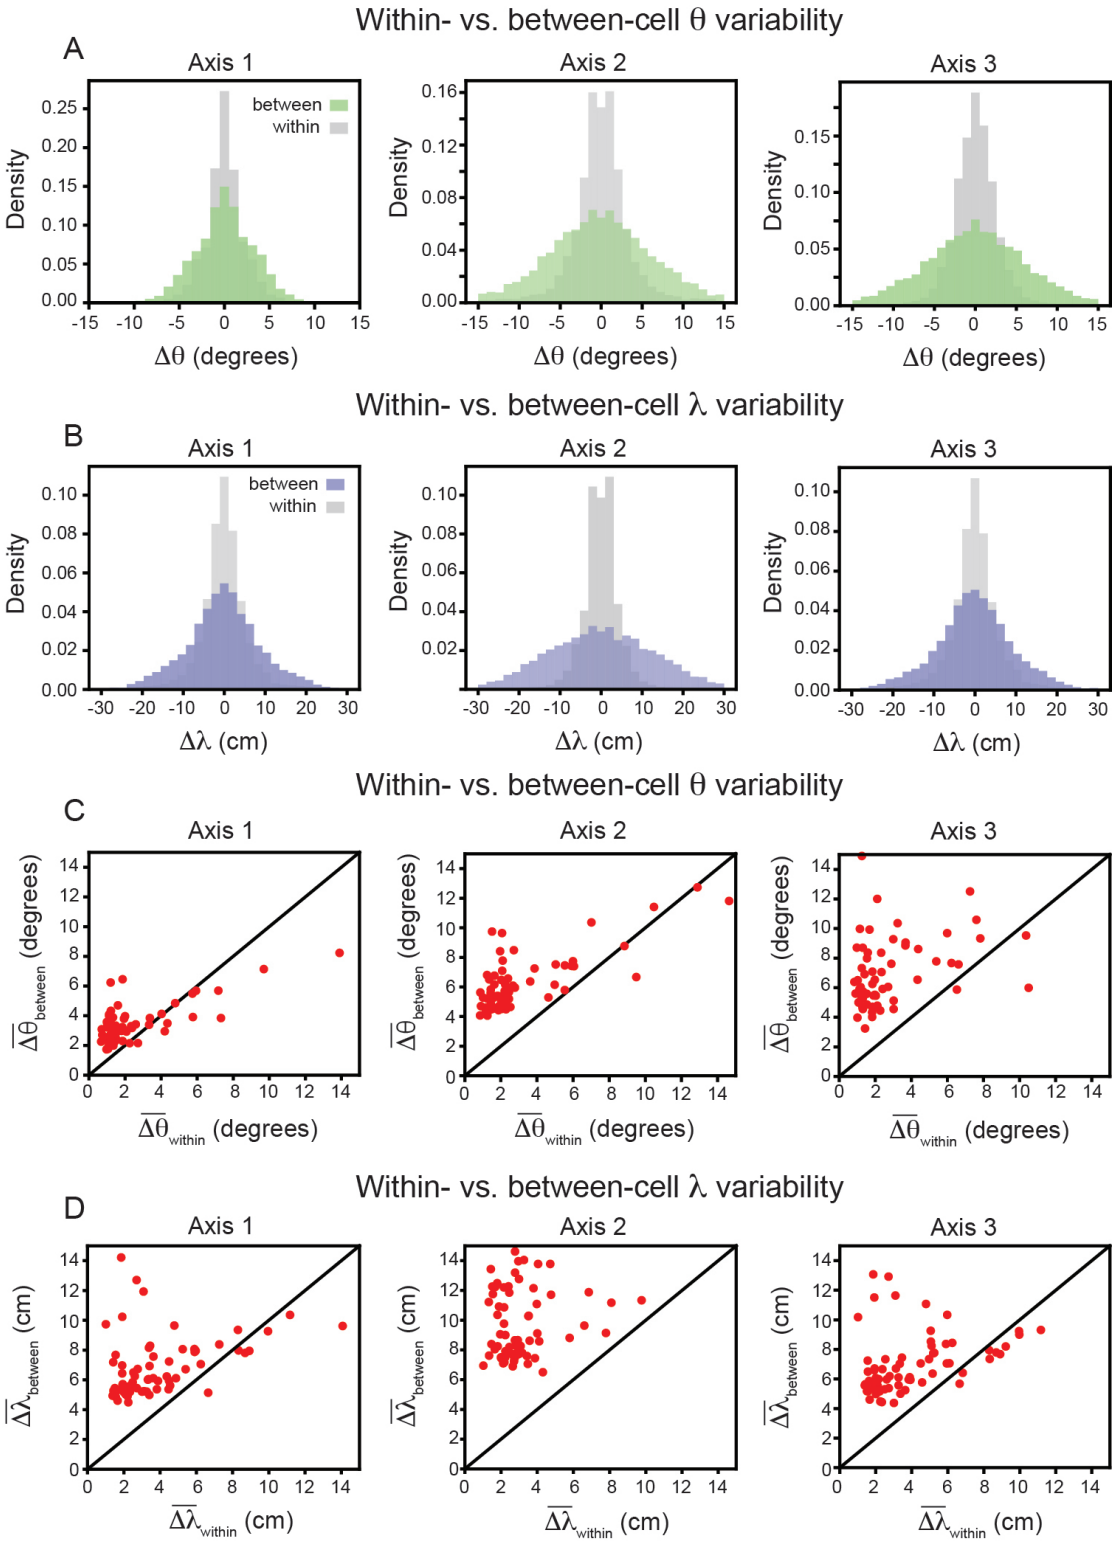

Figure 4: **Variability of grid properties, restricted to the same axis, is a robust property of individual grid module (recording ID R12).** Same analysis as in Fig. 3 (B) – (E), but for variability measured on each axis independently. For visualization, we excluded a small number of cells that were outside the axes limits, including 2, 5, and 10 cells for Axes 1–3, respectively (C); and 3, 4, and 3 cells for Axes 1–3, respectively (D); these cells were included in non-parametric statistical analyses.

no cells that passed our criteria) [34]. Although there are differences across recordings, we find larger between- than within-cell variability for grid orientation and spacing is present across all recordings (Fig. 5A, B – 80.0% of grid cells have greater between- than within-variability for orientation, $\overline{\Delta\theta}_{\text{within}} = 1.4°$ and $\overline{\Delta\theta}_{\text{between}} = 1.9°$; 87.9% of grid cells have greater between- than within-variability for spacing, $\overline{\Delta\lambda}_{\text{within}} = 1.7$ cm and $\overline{\Delta\lambda}_{\text{between}} = 3.6$ cm). A Wilcoxon-Signed-Rank Test finds that these differences are significant ($\Delta\theta$: $p < 0.001$, $\Delta\lambda$:, $p < 0.001$). We do not find evidence suggesting that the within-cell variability is influenced by grid score (Fig. 5C, D – linear regression for θ: $R^2 = 0.03$, $p = 0.55$ Wald Test; linear regression for λ: $R^2 = 0.07$, $p = 0.14$ Wald Test).

Taken together, our analysis of large-scale MEC recordings demonstrates that grid cells in the same grid module do not have a single grid spacing and orientation, but instead have a restricted range of values around the module mean. This property is a robust feature that cannot be explained by noise from the recording or fitting procedures.

Variability in grid properties within single modules improves the encoding of local space

The variability of grid cell properties within individual grid modules, while statistically significant, is small in magnitude. Can a computational benefit in the encoding of local space be gained from this level of inhomogeneity? How sensitive might such a computational benefit be to the exact amount of variability present?

To test this in a setting in which we have complete control over the distribution of grid cell properties, we generated populations of synthetic grid cells (see Materials). For simplicity, we assumed that all grid cells in our population have grid orientation and spacing sampled from Gaussian distributions, with means $\hat{\lambda}$ and $\hat{\theta}$ and standard deviations σ_λ and σ_θ, respectively. Assigning each grid cell in our population a grid orientation and spacing, we are able to generate "ideal" rate maps [14]. Sampling from a Poisson process on these ideal rate maps, we are able to generate noisy grid cell rate maps (Fig. 6A). Leveraging a simple linear decoder (see Materials), we examine how decoding error of local space is affected by the number of grid cells in the population and the amount of variability in λ and θ.

We begin by investigating the decoding capabilities in a synthetic module with properties similar to that of the experimentally recorded module that we analyzed in detail (recording identifier R12 – Figs. 2, 3), setting $\hat{\lambda} = 85$ cm and $\hat{\theta} = 6°$. As expected, when this population of synthetic grid cells has no variability in λ and θ ($\sigma_\lambda = 0$ cm and $\sigma_\theta = 0°$), increasing the number of neurons leads to little change in resolution of local space, as decoding error remains high (Fig. 6D, black line). To determine an appropriate value for σ_λ and σ_θ, we subtract the mean between-cell variability by the mean within-cell variability ($\langle\overline{\Delta\theta}_{\text{between}}\rangle - \langle\overline{\Delta\theta}_{\text{within}}\rangle = 2.0° - 1.2° = 0.8°$ and $\langle\overline{\Delta\lambda}_{\text{between}}\rangle - \langle\overline{\Delta\lambda}_{\text{within}}\rangle = 6.9$ cm $- 1.9$ cm $= 5.0$ cm, where $\langle\cdot\rangle$ denotes mean across all grid cells). We interpret these values as the amount of variability in grid spacing and orientation that is not due to noise. We therefore examine the ability to decode local spatial information in the same synthetic module, now with $\sigma_\theta = 1°$ and $\sigma_\lambda = 5$ cm. This amount of variability leads to similar sampled distributions of λ and θ as was found in the real data (compare Fig. 2B, C with Fig. 6B, C). For this variable population, there is a steady decrease in decoding error with increasing number of grid cells (N), reaching nearly 10 cm decoding error when $N = 1024$ (Fig. 6D).

To determine the extent to which the decrease in decoding error depends the exact variability of grid spacing and orientation, we performed a grid search over 36 pairs of $(\sigma_\lambda, \sigma_\theta)$ values in [0°, 1°, ... , 5°] × [0 cm, 1 cm, ... , 5 cm]. As expected, when the variability of both grid properties was large, we found nearly zero decoding error of local space (Fig. 6E – bottom right). In addition, there is a large range of variability values that lead to improved decoding, including values smaller than used (Fig. 6E – left of white star).

We perform the same grid search on syntethic populations of grid cells with different spacings: $\hat{\lambda} = 50$ cm and $\hat{\lambda} = 100$ cm (Fig. 7), finding it again possible for the decoding error to drop well below that of "fixed" population (i.e., $\sigma_\theta = 0°$ and $\sigma_\lambda = 0$ cm), if a sufficient amount of variability exists. For the synthetic module with smaller spacing ($\hat{\lambda} = 50$ cm), the existence of more grid fields in the 1.5 m × 1.5 m arena leads to greater interference patterns, enabling the same amount of variability in grid orientation to lead

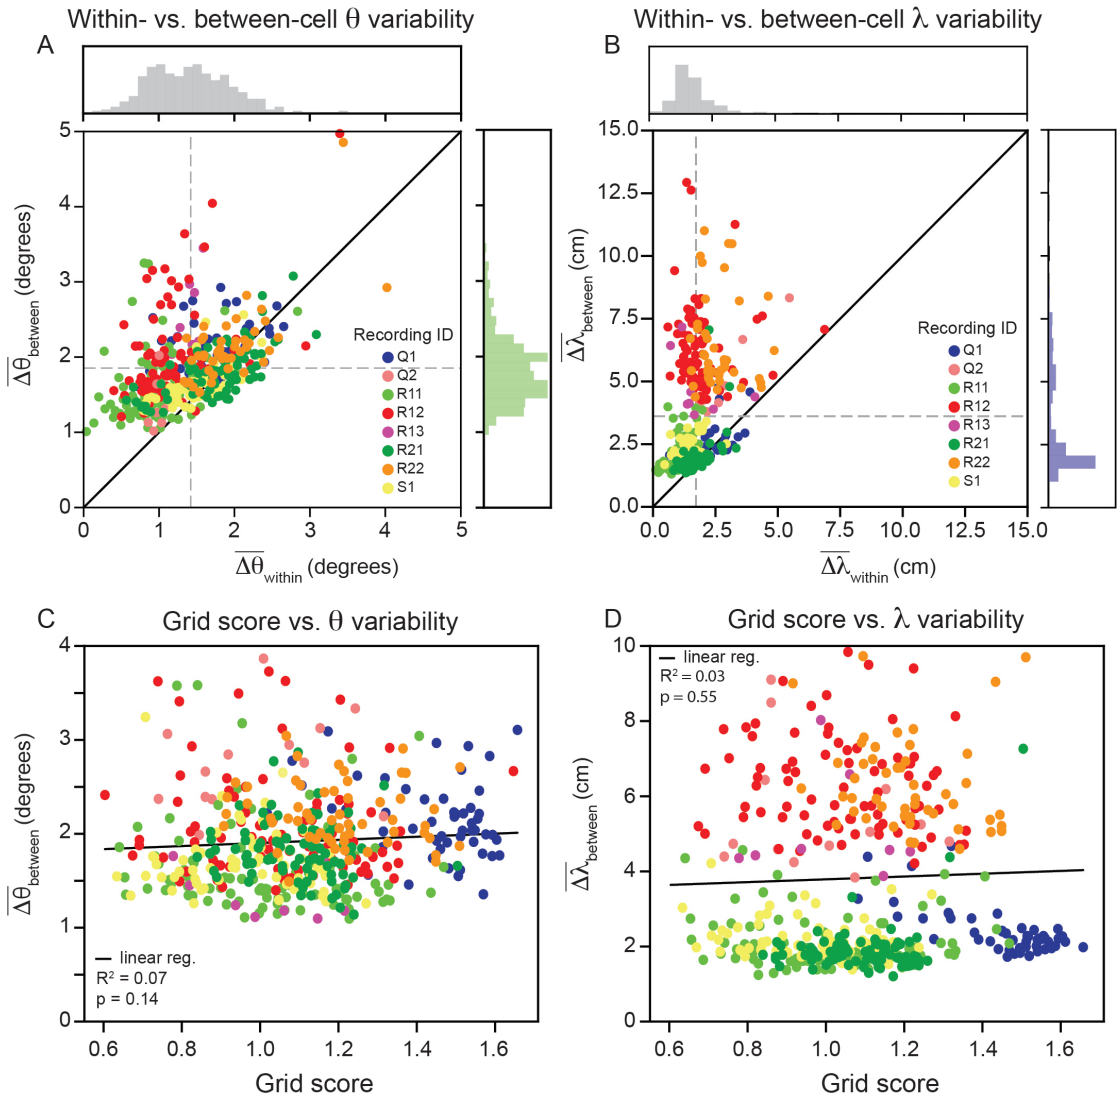

Figure 5: **Within module grid property variability is a robust feature across modules.** (A) Average within-cell variability of grid orientation ($\overline{\Delta\theta}_{\text{within}}$), compared to average between-cell variability of grid orientation ($\overline{\Delta\theta}_{\text{between}}$) for each cell ($N = 420$) across eight recordings (cells colored by their corresponding recording ID). The histogram above the plot shows the distribution of $\overline{\Delta\theta}_{\text{within}}$ and the histogram to the right shows the distribution of $\overline{\Delta\theta}_{\text{between}}$. (B) Same as (A), but for grid spacing. For visualization, 5 cells excluded ($\overline{\Delta\lambda}_{\text{between}} > 15$ cm), but were included in non-parametric statistical analysis. Dashed gray lines show the population mean. (C)–(D) Average within cell variability of θ and λ (respectively), as a function of grid score. For visualization, 3 and 22 cells were excluded for (C)–(D), respectively, but were included in statistical analysis. Black solid line is linear regression, with R^2 and p-value reported.

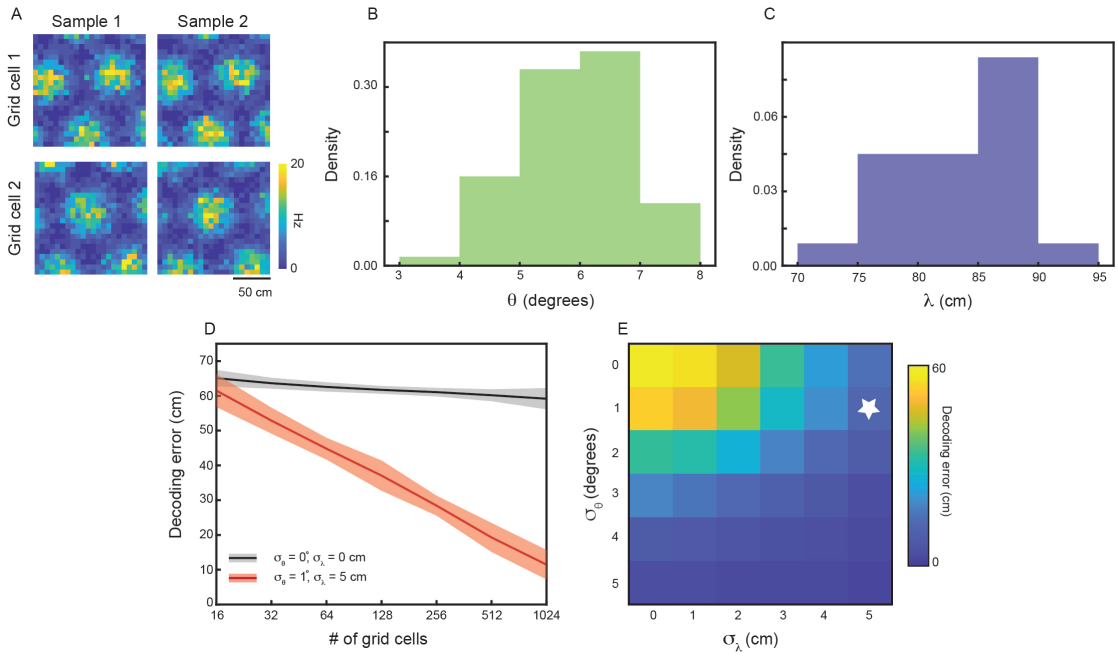

Figure 6: **Variability in grid properties enables improved decoding of local space from the activity of grid cells within a single module.** (A) Example noisy grid cell rate maps generated from Poisson process. The size of the square arena is taken to be 1.5 m $\times$ 1.5 m, to be consistent with what was used in the experimental set-up analyzed [34]. Each row represents a different grid cell, with each column representing a different sampling of the Poisson process. (B)–(C) Distribution of sampled grid spacing and orientation from synthetic population, when using $\sigma_\theta = 1^\circ$ and $\sigma_\lambda = 5$ cm. (D) Decoding error, as a function of grid cell population size, with populations having either no variability in grid properties (black line) or variability like what was present in the data analyzed (red line). Solid line is mean across 25 independent grid cell populations and shaded area is $\pm$ standard deviation of the 25 independent populations. (E) Decoding error, over a grid of σ_θ and σ_λ values, for populations of 1024 grid cells. White star denotes values used in (D). Average grid spacing and orientation taken to be $\hat{\lambda} = 85$ cm and $\hat{\theta} = 6^\circ$, to approximate the results found in recording R12 (Fig. 3) for all subplots.

to a sharper decrease in decoding error. Additionally, for the synthetic module with smaller spacing, a given value of σ_λ is a larger percentage of $\hat{\lambda}$, as compared to the synthetic module with larger spacing. This enables the same amount of variability in grid spacing to lead to a sharper decrease in decoding error, for the synthetic module with the smaller spacing.

Taken together, these results suggest that *individual* grid modules can exhibit significant encoding of local space via heterogeneity in their grid properties, even when the extent of the variability in θ and λ is similar to that found in the analysis of the experimental recordings.

Discussion

The multiple firing fields of the grid cells, organized along a triangular lattice, has been historically interpreted as a limiting feature for encoding of local space, but a powerful feature for enabling path integration [16]. Particularly influential in shaping this view has been the discovery of the distribution of grid cells into distinct modules, with grid cell spacing (λ) and orientation (θ) preserved within, but not across, modules [1, 22], making integration across multiple modules necessary for local spatial information to be decoded [24–28]. While evidence for discontinuity in the grid cell properties across modules is strong, the corollary assumption, that within-module the values of λ and θ are identical (within the bounds of noise), has not been systematically studied.

Analyzing recently collected MEC recordings, we found the range of λ and θ values was large, with examples of grid cell pairs in the same module having over 7° difference in grid orientation and 20 cm

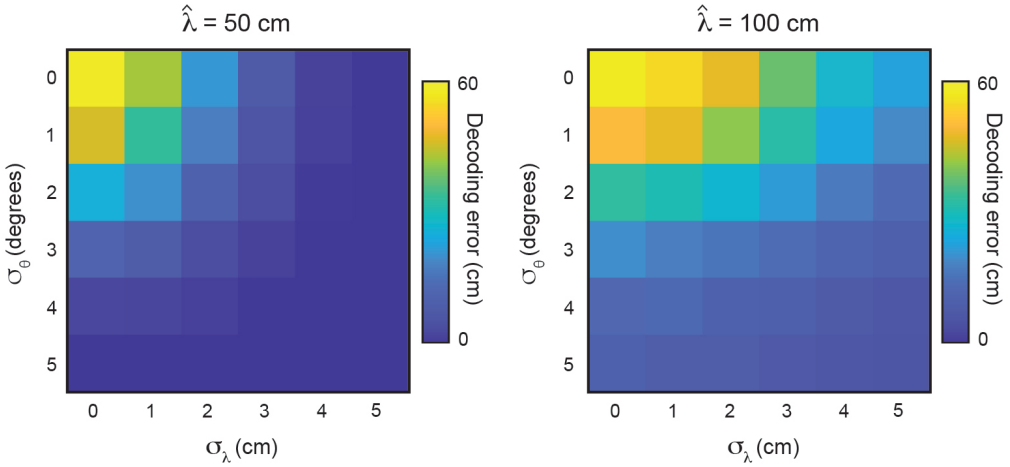

Figure 7: **Variability in grid properties enables improved decoding of local space for grid modules with different grid spacings.** Mean decoding error, over a grid of σ_λ and σ_θ values, for $\hat{\lambda} = 50$ cm and $\hat{\lambda} = 100$ cm, respectively. $\hat{\sigma}$ is set to 0°. Mean is computed across 25 independent grid cell populations of ($N = 1024$ synthetic grid cells).

difference in grid spacing (Fig. 2 E, F). Statistical analysis shows that the variability is more robust than expected from noise for the majority of grid cells (Fig. 5). This was despite the fact that we used a very conservative criteria for assessing whether a grid cell was robust enough to include in the analysis.

Our comparison of within- and between-cell grid property variability was key to our argument, as it was for previous work using it to study the robustness of differences in peak grid cell firing [39]. The absence of its use in the characterization of distribution of λ and θ [30] may be why the heterogeneity was not identified until now. We find that our analysis is robust to choices of parameters (Figs. S1, S2) and whether we treat each grid field independently or take the average across grid fields (Fig. 4). This challenges the commonly held belief that the variability observed in the grid orientation and spacing is attributable solely to noise.

We hypothesize that this variability may be used to increase the fidelity at which *individual* grid modules can encode local space. This idea is consistent with a large body of literature showing that heterogeneity in the responses of populations of neurons increases the robustness of encoding [46–49]. We indeed find, in noisy synthetic models of grid cells, that a level of variability in grid properties similar to what is quantified in the real data leads to more accurate encoding of spatial position (Fig. 6). This benefit is increased with larger numbers of grid cells in the population (Fig. 6D), and is observed over a range of underlying variability values (Fig. 6E). We find that the improvement was most pronounced in modules with small grid spacing (Fig. 7A), although larger modules can see a decrease in decoding error for amounts of variability consistent with was was found in the analyzed data (Fig. 7B).

We note that our results are additionally aligned with recent findings of heterogeneity in maximum firing rates across individual grid fields [37–39] and grid sheering [35, 36, 40]. Our work further demonstrates that, even in the absence of these perturbations, individual grid modules may encode considerably more local spatial information than previously believed.

Finally, models of the formation of orientation maps in visual cortex have demonstrated that slight angular offsets of retinal mosaics, along which retinal receptive fields are organized, can generate complex patterns [50], similar to those found in visual cortex orientation maps [51]. Our results indicate that grid cells in MEC are capable of leveraging a similar computational principle, suggesting that mosaic patterns might be a broadly utilized feature in neural coding [52].

Limitations

While the data set we analyzed [34] represents an advance in the ability to simultaneously record from tens to hundreds of putative grid cells, across grid modules, the MEC remains a challenging brain region to

access for large-scale neurophysiological experiments. Indeed, with our conservative inclusion criteria, we were ultimately limited by having only 420 grid cells included in our analysis. Future work can perform more detailed and complete characterizations of grid property heterogeneity, as new neurotechnologies that enable larger yield of grid cells are developed [53, 54].

Our decoding analysis, while demonstrating the possibility that variability in grid properties can be leveraged by single grid modules to enhance the encoding of local spatial information, made several simplifying assumptions: Poisson sampling for spike rates; linear decoding; normal distribution of grid properties. Future work can assess the extent to which these restrictions can be lifted, while still enabling individual grid modules to have low decoding error.

Open questions

The heterogeneity in grid properties we characterize motivates the investigation of several new lines of research. Because these are directions that we believe to be fruitful for the field as a whole, we outline them below, with our hypotheses for possible answers.

Q1: What causes grid property variability? A natural first direction to address is identifying the source of the heterogeneity in grid cell properties we observe. One hypothesis is that this could be driven by "defects" in the specific connectivity pattern that is needed for a continuous attractor [16]. However, whether stable path integration could occur, in the face of such defects, is unclear. As an alternative, the coupling between hippocampus and MEC, which has been shown to lead to variability in grid field firing rates (as experimentally observed) [39, 55], may lead to differences in grid orientation and spacing. In particular, a previous computational model that learned grid cells from place cell input, using non-negative principal component analysis, has shown that place field width affects grid cell orientation and spacing [6]. Heterogeneity in the spatial coding properties of place cells that has been found along the transverse axis of CA3 [56–58], suggesting there may be a systematic differences in the place field widths of hippocampal inputs to MEC grid cells. Further, it was shown that this place-to-grid cell computational model has a linear relationship between place field width and grid spacing, and a non-monotonic relationship between place field width and grid orientation [6]. These may explain why we find stronger average variability in grid spacing than grid orientation (Fig. 5A, B).

Q2: How does grid property variability affect continuous attractor network structure? Continuous attractor network models of grid cells [2, 3, 16] enjoy considerable experimental support [5, 23, 30, 31, 34], making them one of the "canonical" models in neuroscience. However, these models make use of the assumption that all the grid cells in a given module have the same grid orientation and spacing to simplify the network connectivity. In particular, by assuming equal grid orientation and spacing, it becomes possible to arrange grid cells in a two-dimensional space spanned by their phases (i.e., a "neural sheet" [16]). Neurons close in this space are wired with excitatory connections and neurons far in this space are wired with inhibitory connections. As the data set we analyzed was found to have strong support for the basic predictions of continuous attractor networks (e.g., toroidal topology of the activity manifold) [34], we do not view our results as challenging these models, but consider the question of how these models can accommodate the observed variability an interesting future direction. It is possible that these modifications are learnable, via some plasticity mechanism. These solutions may lead to geometric effects on the underlying manifold, which could be explored by using computational tools [59], as well as the effect the dynamical properties of the neural population activity, which could be quantified by data-driven dynamical systems methods [60–62]. We hypothesize that the degree in variability of grid spacing and orientation may strike a balance between being small enough to keep the continuous attractor network structure stable, but large enough to enable encoding of local information of space.

Q3: How does grid property variability shape hippocampal representations? The projections from MEC to hippocampus suggest that the variability in grid properties may influence hippocampal representations (even if grid cells do not comprise the majority of inputs). We consider two possible ways in which this may happen. First, given that grid cells have been reported to maintain their grid spacing and

orientation across exposures to new environments, while undergoing a change in their grid phase [30, 63], the integration across multiple modules has been necessary to explain place field remapping. However, grid phase plays an important role in generating the specific complex interference patterns that emerge when considering the joint activity of grid cells with variable grid properties (Fig. 1), and thus the reported changes in phase may be sufficient to generate large (and seemingly random) changes in local spatial information conveyed by grid cells to hippocampal cells. This could drive additional changes in the local spatial information projected to hippocampus, as well as explain the significant differences in correlation structure between CA1 neurons across different environments [64].

Recent work on hippocampal place field drift [65–69] has demonstrated that there is a significant change in the place field location across time, especially in CA1. One possible source of this phenomenon is the reported instability of dendritic spines on the apical dendrites of CA1 place cells [58, 70–72], ostensibly leading to changes in the MEC inputs to these neurons. However, if grid cells across multiple modules are necessary for local spatial information, turnover in synaptic input is unlikely to cause large changes in the spatial preferences of CA1 neurons, as integration over several scales should provide stable encoding properties. In contrast, different subpopulations of grid cells with variable grid properties can lead to differences in the local spatial information encoded by their joint activity, even if they come from a single module, possibly influencing the spatial preferences of CA1 place cells.

Q4: Why are there multiple modules? Given that our results demonstrate the ability of single grid modules to encode information about local space – a feat previously believed to be possible only if activity from multiple grid modules was integrated together – why is the MEC organized into multiple modules? While the variability removes the necessity for encoding local space with multiple modules, higher fidelity representations may be achievable by integrating across multiple modules. In addition, the use of grid cells beyond spatial navigation [73–78], where hierarchical representations are important [42, 43], may be a sufficient implicit bias for the formation of multiple modules [79]. In particular, encoding information at multiple distinct scales is critical for multi-scale reasoning, a cognitive function grid cells may support.

Materials

Electrophysiology recordings

The neural activity analyzed in this paper comes from a publicly available data set, which has previously been described in detail [34]. We provide brief summary of the methodology and the experimental paradigms used during the recordings.

Three male rats (Long Evans – Rats Q, R, and S) were implanted with Neuropixels silicon probes [80, 81]. These probes were targeted at the MEC-parasubiculum region and the surgery was performed as described previously [32, 81]. After three hours of recovery, recordings were performed.

In the recordings analyzed, the rats foraged for randomly dispersed corn puffs in a 1.5×1.5 m^2 square open field arena, with walls of height 50 cm. The rats were familiar with the environment and task, having trained 10–20 times prior to the implantation. The rats were food restricted to motivate their foraging, being kept at a minimum of 90% of their original body weight (300–500 grams).

All procedures in the original study were approved by the Norwegian Food and Safety Authority and done in accordance with the Norwegian Animal Welfare Act and the European Convention for the Protection of Vertebrate Animals used for Experimental and Other Scientific Purposes.

Electrophysiology post-processing

The neural activity analyzed in this paper was post-processed, before made publicly available. We describe, in brief, the post-processing performed [34], as well as the post-processing we performed on the downloaded data.

Spike sorting, via KiloSort 2.5 [81], was applied to the data recorded from the Neuropixel probes. Individual units were deemed putative cells if their average spike rate was in the range of 0.5–10Hz, and 99% of their interspike intervals were greater than 2 ms.

For each putative cell, rate maps were constructed by averaging the activity at binned spatial positions in the open field arena. This raw rate map was smoothed, using a Gaussian kernel. The autocorrelation of these rate maps were computed, and a grid score calculated, as described previously [44]. To determine which putative cells were grid cells, and to which module they belonged, autocorrelograms of raw rate maps were constructed, with some of the bins corresponding to the closest and furthest distances removed. These "cropped" rate maps were then clustered by using UMAP [82] to project the data onto a two-dimensional subspace. All the cells in the non-largest cluster were found to have similar grid spacing and orientation, as well as high grid scores. They were therefore deemed putative grid cells. All cells in these clusters had their spike trains placed in the public repository from which we downloaded the data.

From the downloaded spike trains, we constructed rate maps and autocorrelograms in a similar manner, using code made publicly available [8] (Fig. 2).

Computing grid score orientation and spacing

Grid spacing and grid orientation were computed according to standard methods described in detail previously [22, 83]. Briefly, the goal of the procedure is to identify the location of the six nearest fields in the spatial autocorrelogram (SAC) (Fig. 2A). This was achieved by performing the following steps. First, the SAC was smoothed using a 2D Gaussian filter with $\sigma = 1$. Second, the center of the SAC was excluded by applying a mask (i.e., a matrix whose center elements were 0 and all other elements were 1). Because the size of the SAC's central peak changes for every module, the radius of such mask (determining the boundary between 1's and 0's) was re-computed for each module. Third, we thresholded the SAC by applying an extended-maxima transform `ndimage.maximum_filter`. Fourth, we identified the center of every field by using the function `scipy.stats.find_peaks`.

Once the peaks of every field had been found, we computed the location of every peak in polar coordinates. We then selected the 6 peaks that were closest to the center of the SAC, based on the computed radial components. Because every SAC is symmetric, we considered for further analysis the 3 peaks closest to the X axis in angular distance (Axis 1, 2, and 3 [45]). Grid spacing was computed as the arithmetic mean of the radial component of the 3 peaks (except when each peak was analyzed separately – Fig. 4). Given that the SAC dimensions are twice of that of the real arena, we multiplied the SAC radial mean by a factor of 2. Grid orientation was computed as the angular mean of orientations (relative to the x-axis) of the three peaks (except when each peak was analyzed separately – Fig. 4).

In order to ensure that subsequent analysis was performed only on cells whose SAC's could be well described by hexagonal structure, we imposed the following constraints: 1) the relative angle between two peaks could not be $< 30°$; 2) the relative angle between two peaks could not be $> 90°$; 3) the values for grid spacing between the peaks could not be substantially different (the ratio of spacings between any two peaks must be > 0.5 and < 2). Cells that did not meet these criteria were determined to have "Poor grid fit" and were rejected from all subsequent analysis. The percentage of all cells that were removed by this inclusion criteria is shown in Fig. S1. We note that cells from one of the nine modules in the publicly available data set that we analyzed had 98.4% of cells rejected by this criteria (recording identifier – R23). We therefore did not include it in any subsequent analysis.

To compute the grid score of recorded MEC cells, we made use of previously published code [8], that is based on metrics that have become standards in quantifying grid cell properties[44].

For analysis of the distribution of grid properties (Fig. 2E, F), we included only grid cells with grid scores greater than 0.85. This was done to demonstrate that variability was present even in cells that exhibit robust grid cell properties. In the subsequent analyses, an alternative criteria is used, which considers the reliability of the grid responses (see below).

Within and between cells splits

To characterize the within- and between-cell variability of grid spacing and orientation, we employed the following approach. First, we split the data into bins of fixed length (30 seconds). From this, we randomly assigned each interval to one of two blocks (denoted as blocks A and B), with exactly half the total number of intervals in each block. For each grid cell, we computed the grid spacing [$\lambda_A^{(i)}$ and $\lambda_B^{(i)}$] and orientation [$\theta_A^{(i)}$ and $\theta_B^{(i)}$], from the data in each block. The within-cell differences of grid spacing and orientation was determined as $\Delta\theta_{\text{within}}^{(i)} = \theta_A^{(i)} - \theta_B^{(i)}$ and $\Delta\lambda_{\text{within}}^{(i)} = \lambda_A^{(i)} - \lambda_B^{(i)}$. Each grid cell's properties were also compared those of another grid cell, with the match being made using random sampling without replacement. These comparisons were determined as $\Delta\theta_{\text{between}}^{(i,j)} = \theta_A^{(i)} - \theta_B^{(j)}$ and $\Delta\lambda_{\text{between}}^{(i,j)} = \lambda_A^{(i)} - \lambda_B^{(j)}$. This process was repeated 100 times (referring to each iteration as a "shuffle"), per recording. Examples of splits of the data, for different shuffles, is schematically illustrated in Fig. 3A). The resulting distributions of $\Delta\theta_{\text{within}}$, $\Delta\lambda_{\text{within}}$, $\Delta\theta_{\text{between}}$, and $\Delta\lambda_{\text{between}}$, across grid cells and splits of the data were then compared. If the within-cell variability in grid spacing and orientation was smaller than between-cell variability, we concluded that the variability of grid cell properties was a robust feature of the data and not due to noise.

When performing this shuffle analysis, we found that some cells, despite having a good grid fit when all the data was considered, did not have SACs that were well described by hexagonal structure when the data was split in half. We viewed this a manifestation of unreliable grid coding and a possible confound in our quantification of variability. As such, we introduced a new inclusion criteria (replacing that of requiring a grid score of > 0.85), only considering cells that had poor grid fits on $< 5\%$ of all shuffles. The percentage of all cells that were removed by this inclusion criteria is shown in Fig. S1. In general, we found that cells with high grid score were reliable, although there were exceptions. Additionally, we found that size of the bin used for splitting the data did not significantly affect the percent of cells with good grid fits (passed the prior inclusion criteria) that were considered reliable (Fig. S2).

Synthetic grid cells

To study how variability in grid cell properties might endow the grid code with computational advantages, we generated synthetic grid cell rate maps, so that we could have complete control over the distribution of their properties. These synthetic grid cell rate maps were constructed as follows.

First, the lengths of each dimension the "arena" within which the simulated grid cells exist (L_x and L_y) were set. Then, for $N \in \mathbb{N}$ grid cells, the grid spacing and orientation were sampled via $\lambda^{(i)} \sim \mathcal{N}(\mu_\lambda, \sigma_\lambda)$ and $\theta^{(i)} \sim \mathcal{N}(\mu_\theta, \sigma_\theta)$ where $\mathcal{N}(\mu, \sigma)$ is a normal distribution with mean μ and variance σ^2. Grid phase was sampled as $\phi^{(i)} \sim \mathcal{U}([0, L_x] \times [0, L_y])$, where ϕ is a two dimensional vector, with first component uniformly sampled from $[0, L_x]$ and second component uniformly sampled from $[0, L_y]$. To construct a population with no variability in grid properties (to use as a control), we set $\sigma_\lambda = 0$ m. and $\sigma_\theta = 0°$.

For each grid cell, we generated idealized grid responses by summing three two-dimensional sinusoids [14], such that the activity at $\mathbf{x} = (x, y) \in [-L_x/2, L_x/2] \times [-L_y/2, L_y/2]$ is given by

$$X^{(i)}(\mathbf{x}) = X^{\max}\frac{2}{3}\left(\frac{1}{3}\sum_{j=1}^{3}\cos\left[R(\theta_i)\mathbf{k}_j(\mathbf{x}+\phi_i)\pi/\lambda_i\right]\right), \tag{1}$$

where $X^{\max}$ is the maximal firing rate, $R(\theta_i)$ is the two-dimensional rotation matrix, with rotation θ_i, and $\mathbf{k}_j$ are the wave vectors with 0°, 60°, and 120° angular differences [i.e., $\mathbf{k}_1 = (1, 0)^T, \mathbf{k}_2 = (1/2, \sqrt{3}/2)^T, \mathbf{k}_3 = (1/2, -\sqrt{3}/2)^T$].

To determine the extent to which local spatial information can be decoded from the activity of populations of grid cells with different degrees of variability in their grid properties, we performed the following analysis.

We generated noisy synthetic spike rates of N grid cells by assuming a Poisson process and sampling using the idealized rate maps (Eq. 1). More concretely, the activity of grid cell i at position (x, y) was assumed to be a random variable with a Poisson distribution, whose mean was $X^{(i)}(x, y)$. Thus, the probability of

observing $\tilde{X}^{(i)}(x,y)$ spikes, at position (x,y), is given by

$$\tilde{X}^{(i)}(x,y) \sim \mathcal{P}\left(X^{(i)}(x,y)\right). \tag{2}$$

Linear decoding

For a given resolution of the arena, we generated $\tilde{X}_j^{(i)}(x,y)$, where $j = 1, ..., 10$. That is, we constructed 10 noisy rate maps. We performed cross-validated decoding by averaging across 9 of the 10 rate maps, to get an average rate map $\hat{X}^{(i)}(x,y)$. For sake of simplicity, consider $\hat{X}^{(i)}(x,y) = 1/9 \sum_{j=1}^{9} \tilde{X}_j^{(i)}(x,y)$. To decode local position from the held out noisy rate map [e.g. $\tilde{X}_{10}^{(i)}(x,y)$], we multiplied the activity at each position by all the positions in the average rate map, taking the sum and assigning the decoded position as that with the largest value. The Euclidean distance between the decoded position and the true position is considered the error. This was performed 10 times (holding out each rate map once) and the average error across all positions in the environment was then averaged across all 10 of the validation splits.

Acknowledgments

We thank the members of the Goard Lab, Francisco Acosta, Spencer Smith, and the DYNS graduate students for useful discussions surrounding this work.

References

[1] Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.

[2] Alexis Guanella and Paul FMJ Verschure. A model of grid cells based on a path integration mechanism. In *Artificial Neural Networks–ICANN 2006: 16th International Conference, Athens, Greece, September 10-14, 2006. Proceedings, Part I 16*, pages 740–749. Springer, 2006.

[3] Mark C Fuhs and David S Touretzky. A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*, 26(16):4266–4276, 2006.

[4] Hugh T Blair, Adam C Welday, and Kechen Zhang. Scale-invariant memory representations emerge from moire interference between grid fields that produce theta oscillations: a computational model. *Journal of Neuroscience*, 27(12):3211–3229, 2007.

[5] Jonathan J Couey, Aree Witoelar, Sheng-Jia Zhang, Kang Zheng, Jing Ye, Benjamin Dunn, Rafal Czajkowski, May-Britt Moser, Edvard I Moser, Yasser Roudi, et al. Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature neuroscience*, 16(3):318–324, 2013.

[6] Yedidyah Dordek, Daniel Soudry, Ron Meir, and Dori Derdikman. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *Elife*, 5:e10094, 2016.

[7] Christopher J. Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In *International Conference on Learning Representations*, 2018.

[8] Andrea Banino, Caswell Barry, Benigno Uria, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J Chadwick, Thomas Degris, Joseph Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, 2018.

[9] Simon Nikolaus Weber and Henning Sprekeler. Learning place cells, grid cells and invariances with excitatory and inhibitory plasticity. *Elife*, 7:e34560, 2018.

[10] Ben Sorscher, Gabriel Mel, Surya Ganguli, and Samuel Ocko. A unified theory for the origin of grid cells through the lens of pattern formation. *Advances in neural information processing systems*, 32, 2019.

[11] Mikail Khona, Sarthak Chandra, and Ila R Fiete. From smooth cortical gradients to discrete modules: spontaneous and topologically robust emergence of modularity in grid cells. *bioRxiv*, pages 2021–10, 2022.

[12] Ben Sorscher, Gabriel C Mel, Samuel A Ocko, Lisa M Giocomo, and Surya Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 2022.

[13] Bruce L McNaughton, Francesco P Battaglia, Ole Jensen, Edvard I Moser, and May-Britt Moser. Path integration and the neural basis of the'cognitive map'. *Nature Reviews Neuroscience*, 7(8):663–678, 2006.

[14] Trygve Solstad, Edvard I Moser, and Gaute T Einevoll. From grid cells to place cells: a mathematical model. *Hippocampus*, 16(12):1026–1031, 2006.

[15] Edmund T Rolls, Simon M Stringer, and Thomas Elliot. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Network: Computation in Neural Systems*, 17(4):447–465, 2006.

[16] Yoram Burak and Ila R Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS computational biology*, 5(2):e1000291, 2009.

[17] Licurgo de Almeida, Marco Idiart, and John E Lisman. The input–output transformation of the hippocampal granule cells: from grid cells to place fields. *Journal of Neuroscience*, 29(23):7504–7512, 2009.

[18] John L Kubie and Steven E Fox. Do the spatial frequencies of grid cells mold the firing fields of place cells? *Proceedings of the National Academy of Sciences*, 112(13):3860–3861, 2015.

[19] Daniel Bush, Caswell Barry, Daniel Manson, and Neil Burgess. Using grid cells for navigation. *Neuron*, 87(3):507–520, 2015.

[20] Jake Ormond and Bruce L McNaughton. Place field expansion after focal mec inactivations is consistent with loss of fourier components and path integrator gain reduction. *Proceedings of the National Academy of Sciences*, 112(13):4116–4121, 2015.

[21] Caitlin S Mallory, Kiah Hardcastle, Jason S Bant, and Lisa M Giocomo. Grid scale drives the scale and long-term stability of place maps. *Nature neuroscience*, 21(2):270–282, 2018.

[22] Hanne Stensola, Tor Stensola, Trygve Solstad, Kristian Frøland, May-Britt Moser, and Edvard I Moser. The entorhinal grid map is discretized. *Nature*, 492(7427):72–78, 2012.

[23] Yi Gu, Sam Lewallen, Amina A Kinkhabwala, Cristina Domnisoru, Kijung Yoon, Jeffrey L Gauthier, Ila R Fiete, and David W Tank. A map-like micro-organization of grid cells in the medial entorhinal cortex. *Cell*, 175(3):736–750, 2018.

[24] Xue-Xin Wei, Jason Prentice, and Vijay Balasubramanian. A principle of economy predicts the functional architecture of grid cells. *Elife*, 4:e08362, 2015.

[25] Martin Stemmler, Alexander Mathis, and Andreas VM Herz. Connecting multiple spatial scales to decode the population activity of grid cells. *Science Advances*, 1(11):e1500816, 2015.

[26] Ila R Fiete, Yoram Burak, and Ted Brookings. What grid cells convey about rat location. *Journal of Neuroscience*, 28(27):6858–6871, 2008.

[27] Sameet Sreenivasan and Ila Fiete. Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature neuroscience*, 14(10):1330–1337, 2011.

[28] Alexander Mathis, Andreas VM Herz, and Martin Stemmler. Optimal population codes for space: grid cells outperform place cells. *Neural computation*, 24(9):2280–2317, 2012.

[29] NM Van Strien, NLM Cappaert, and MP Witter. The anatomy of memory: an interactive overview of the parahippocampal–hippocampal network. *Nature reviews neuroscience*, 10(4):272–282, 2009.

[30] KiJung Yoon, Michael A Buice, Caswell Barry, Robin Hayman, Neil Burgess, and Ila R Fiete. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature neuroscience*, 16(8):1077–1084, 2013.

[31] Benjamin Dunn, Maria Mørreaunet, and Yasser Roudi. Correlations and functional connections in a population of grid cells. *PLoS computational biology*, 11(2):e1004052, 2015.

[32] Richard J Gardner, Li Lu, Tanja Wernle, May-Britt Moser, and Edvard I Moser. Correlation structure of grid cells is preserved during sleep. *Nature neuroscience*, 22(4):598–608, 2019.

[33] Sean G Trettel, John B Trimper, Ernie Hwaun, Ila R Fiete, and Laura Lee Colgin. Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nature neuroscience*, 22(4):609–617, 2019.

[34] Richard J Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A Baas, Benjamin A Dunn, May-Britt Moser, and Edvard I Moser. Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128, 2022.

[35] Dori Derdikman, Jonathan R Whitlock, Albert Tsao, Marianne Fyhn, Torkel Hafting, May-Britt Moser, and Edvard I Moser. Fragmentation of grid cell maps in a multicompartment environment. *Nature neuroscience*, 12(10):1325–1332, 2009.

[36] Julija Krupic, Marius Bauza, Stephen Burton, Caswell Barry, and John O'Keefe. Grid cell symmetry is shaped by environmental geometry. *Nature*, 518(7538):232–235, 2015.

[37] Geoffrey W Diehl, Olivia J Hon, Stefan Leutgeb, and Jill K Leutgeb. Grid and nongrid cells in medial entorhinal cortex represent spatial location and environmental features with complementary coding schemes. *Neuron*, 94(1):83–92, 2017.

[38] Revekka Ismakov, Omri Barak, Kate Jeffery, and Dori Derdikman. Grid cells encode local positional information. *Current Biology*, 27(15):2337–2343, 2017.

[39] Benjamin Dunn, Daniel Wennberg, Ziwei Huang, and Yasser Roudi. Grid cells show field-to-field variability and this explains the aperiodic response of inhibitory interneurons. *arXiv preprint arXiv:1701.04893*, 2017.

[40] Gily Ginosar, Johnatan Aljadeff, Liora Las, Dori Derdikman, and Nachum Ulanovsky. Are grid cells used for navigation? on local metrics, subjective spaces, and black holes. *Neuron*, 2023.

[41] Jeff Hawkins, Marcus Lewis, Mirko Klukas, Scott Purdy, and Subutai Ahmad. A framework for intelligence and cortical function based on grid cells in the neocortex. *Frontiers in neural circuits*, 12:121, 2019.

[42] Mirko Klukas, Marcus Lewis, and Ila Fiete. Efficient and flexible representation of higher-dimensional cognitive variables with grid cells. *PLoS computational biology*, 16(4):e1007796, 2020.

[43] Jon W Rueckemann, Marielena Sosa, Lisa M Giocomo, and Elizabeth A Buffalo. The grid code for ordered experience. *Nature Reviews Neuroscience*, 22(10):637–649, 2021.

[44] Francesca Sargolini, Marianne Fyhn, Torkel Hafting, Bruce L McNaughton, Menno P Witter, May-Britt Moser, and Edvard I Moser. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762, 2006.

[45] Tor Stensola, Hanne Stensola, May-Britt Moser, and Edvard I Moser. Shearing-induced asymmetry in entorhinal grid cells. *Nature*, 518(7538):207–212, 2015.

[46] Maoz Shamir and Haim Sompolinsky. Implications of neuronal diversity on population coding. *Neural computation*, 18(8):1951–1986, 2006.

[47] Mircea I Chelaru and Valentin Dragoi. Efficient coding in heterogeneous neuronal populations. *Proceedings of the National Academy of Sciences*, 105(42):16344–16349, 2008.

[48] Julijana Gjorgjieva, Guillaume Drion, and Eve Marder. Computational implications of biophysical diversity and multiple timescales in neurons and synapses for circuit performance. *Current opinion in neurobiology*, 37:44–52, 2016.

[49] Nicolas Perez-Nieves, Vincent CH Leung, Pier Luigi Dragotti, and Dan FM Goodman. Neural heterogeneity promotes robust learning. *Nature communications*, 12(1):5791, 2021.

[50] Se-Bum Paik and Dario L Ringach. Retinal origin of orientation maps in visual cortex. *Nature neuroscience*, 14(7):919–925, 2011.

[51] Gary G Blasdel and Guy Salama. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321(6070):579–585, 1986.

[52] Spencer L Smith and Ikuko T Smith. life imitates op art. *Nature Neuroscience*, 14(7):803–804, 2011.

[53] Ryan J Low, Yi Gu, and David W Tank. Cellular resolution optical access to brain regions in fissures: imaging medial prefrontal cortex and grid cells in entorhinal cortex. *Proceedings of the National Academy of Sciences*, 111(52):18739–18744, 2014.

[54] Weijian Zong, Horst A Obenhaus, Emilie R Skytøen, Hanna Eneqvist, Nienke L de Jong, Ruben Vale, Marina R Jorge, May-Britt Moser, and Edvard I Moser. Large-scale two-photon calcium imaging in freely moving mice. *Cell*, 185(7):1240–1256, 2022.

[55] Haggai Agmon and Yoram Burak. A theory of joint attractor dynamics in the hippocampus and the entorhinal cortex accounts for artificial remapping and grid cell field-to-field variability. *Elife*, 9:e56894, 2020.

[56] Heekyung Lee, Cheng Wang, Sachin S Deshmukh, and James J Knierim. Neural population evidence of functional heterogeneity along the ca3 transverse axis: pattern completion versus pattern separation. *Neuron*, 87(5):1093–1105, 2015.

[57] Li Lu, Kei M Igarashi, Menno P Witter, Edvard I Moser, and May-Britt Moser. Topography of place maps along the ca3-to-ca2 axis of the hippocampus. *Neuron*, 87(5):1078–1092, 2015.

[58] William T Redman, Nora S Wolcott, Luca Montelisciani, Gabriel Luna, Tyler D Marks, Kevin K Sit, Che-Hang Yu, Spencer Smith, and Michael J Goard. Long-term transverse imaging of the hippocampus with glass microperiscopes. *Elife*, 11:e75391, 2022.

[59] Francisco Acosta, Sophia Sanborn, Khanh Dao Duc, Manu Madhav, and Nina Miolane. Quantifying extrinsic curvature in neural manifolds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 610–619, 2023.

[60] William T Redman, Maria Fonoberova, Ryan Mohr, Ioannis G Kevrekidis, and Igor Mezić. Algorithmic (semi-) conjugacy via koopman operator theory. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 6006–6011. IEEE, 2022.

[61] William T Redman, Juan M Bello-Rivas, Maria Fonoberova, Ryan Mohr, Ioannis G Kevrekidis, and Igor Mezić. On equivalent optimization of machine learning methods. *arXiv preprint arXiv:2302.09160*, 2023.

[62] Mitchell Ostrow, Adam Eisen, Leo Kozachkov, and Ila Fiete. Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis. *Advances in Neural Information Processing Systems*, 36, 2024.

[63] Marianne Fyhn, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I Moser. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–194, 2007.

[64] Eliott Robert Joseph Levy, Simón Carrillo-Segura, Eun Hye Park, William Thomas Redman, José Rafael Hurtado, SueYeon Chung, and André Antonio Fenton. A manifold neural population code for space in hippocampal coactivity dynamics independent of place fields. *Cell Reports*, 42(10), 2023.

[65] Emily A Mankin, Fraser T Sparks, Begum Slayyeh, Robert J Sutherland, Stefan Leutgeb, and Jill K Leutgeb. Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences*, 109(47):19462–19467, 2012.

[66] Yaniv Ziv, Laurie D Burns, Eric D Cocker, Elizabeth O Hamel, Kunal K Ghosh, Lacey J Kitch, Abbas El Gamal, and Mark J Schnitzer. Long-term dynamics of ca1 hippocampal place codes. *Nature neuroscience*, 16(3):264–266, 2013.

[67] Thomas Hainmueller and Marlene Bartos. Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*, 558(7709):292–296, 2018.

[68] Walter G Gonzalez, Hanwen Zhang, Anna Harutyunyan, and Carlos Lois. Persistence of neuronal representations through time and damage in the hippocampus. *Science*, 365(6455):821–825, 2019.

[69] Can Dong, Antoine D Madar, and Mark EJ Sheffield. Distinct place cell dynamics in ca1 and ca3 encode experience in new environments. *Nature communications*, 12(1):2977, 2021.

[70] Adi Mizrahi, Justin C Crowley, Eran Shtoyerman, and Lawrence C Katz. High-resolution in vivo imaging of hippocampal dendrites and spines. *Journal of Neuroscience*, 24(13):3147–3151, 2004.

[71] Alessio Attardo, James E Fitzgerald, and Mark J Schnitzer. Impermanence of dendritic spines in live adult ca1 hippocampus. *Nature*, 523(7562):592–596, 2015.

[72] Thomas Pfeiffer, Stefanie Poll, Stephane Bancelin, Julie Angibaud, VVG Krishna Inavalli, Kevin Keppler, Manuel Mittag, Martin Fuhrmann, and U Valentin Nägerl. Chronic 2p-sted imaging reveals high turnover of dendritic spines in the hippocampus in vivo. *Elife*, 7:e34700, 2018.

[73] Nathaniel J Killian, Michael J Jutras, and Elizabeth A Buffalo. A map of visual space in the primate entorhinal cortex. *Nature*, 491(7426):761–764, 2012.

[74] Alexandra O Constantinescu, Jill X O'Reilly, and Timothy EJ Behrens. Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292):1464–1468, 2016.

[75] Dmitriy Aronov, Rhino Nevers, and David W Tank. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, 543(7647):719–722, 2017.

[76] Joshua B Julian, Alexandra T Keinath, Giulia Frazzetta, and Russell A Epstein. Human entorhinal cortex represents visual space using a boundary-anchored grid. *Nature neuroscience*, 21(2):191–194, 2018.

[77] Matthias Nau, Tobias Navarro Schröder, Jacob LS Bellmund, and Christian F Doeller. Hexadirectional coding of visual space in human entorhinal cortex. *Nature neuroscience*, 21(2):188–190, 2018.

[78] Niklas Wilming, Peter König, Seth König, and Elizabeth A Buffalo. Entorhinal cortex receptive fields are modulated by spatial attention, even without movement. *Elife*, 7:e31745, 2018.

[79] Rylan Schaeffer, Mikail Khona, Tzuhsuan Ma, Cristobal Eyzaguirre, Sanmi Koyejo, and Ila Fiete. Self-supervised learning of representations for space generates multi-modular grid cells. *Advances in Neural Information Processing Systems*, 36, 2024.

[80] James J Jun, Nicholas A Steinmetz, Joshua H Siegle, Daniel J Denman, Marius Bauza, Brian Barbarits, Albert K Lee, Costas A Anastassiou, Alexandru Andrei, Çağatay Aydın, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017.

[81] Nicholas A Steinmetz, Cagatay Aydin, Anna Lebedeva, Michael Okun, Marius Pachitariu, Marius Bauza, Maxime Beau, Jai Bhagat, Claudia Böhm, Martijn Broux, et al. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372(6539):eabf4588, 2021.

[82] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

[83] William N Butler, Kiah Hardcastle, and Lisa M Giocomo. Remembered reward locations restructure entorhinal spatial maps. *Science*, 363(6434):1447–1452, 2019.

Supplemental material

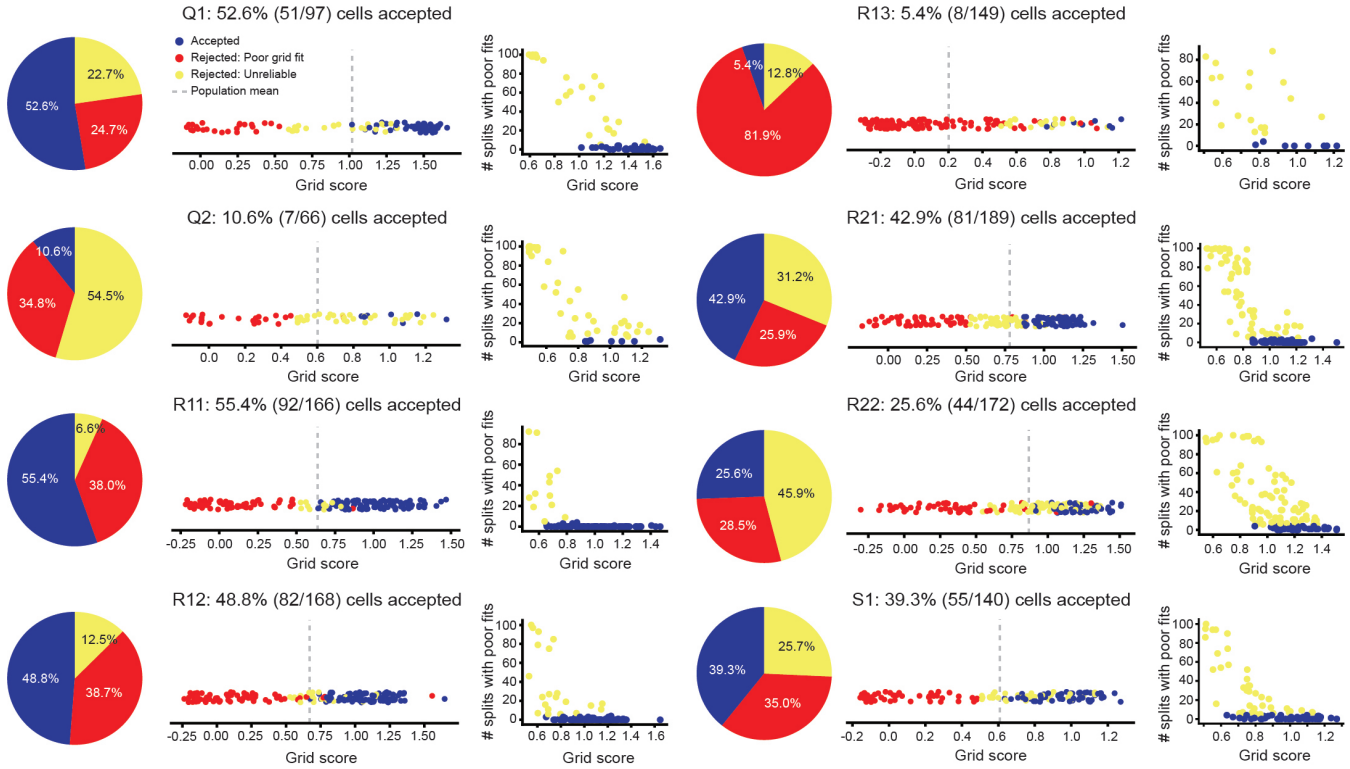

Figure S1: **Accepted and rejected cells in all grid modules.** (Left) The percent of all cells, across all modules, that were rejected by our inclusion criteria. Cells that were rejected as not having SAC's that were well described by hexagonal structure ("Rejected: Poor grid fits") are shown in red and cells that were rejected as not reliably having SAC's that were well described by hexagonal structure ("Rejected: Unreliable") are shown in yellow. Cells that met these criteria ("Accepted") are shown in blue. (Middle) Grid score of all cells, with coloring denoting whether they were accepted or rejected. (Right) The number of splits of the data (out of 100) that cells had SAC's with poor grid fits, as a function of each cell's grid score.

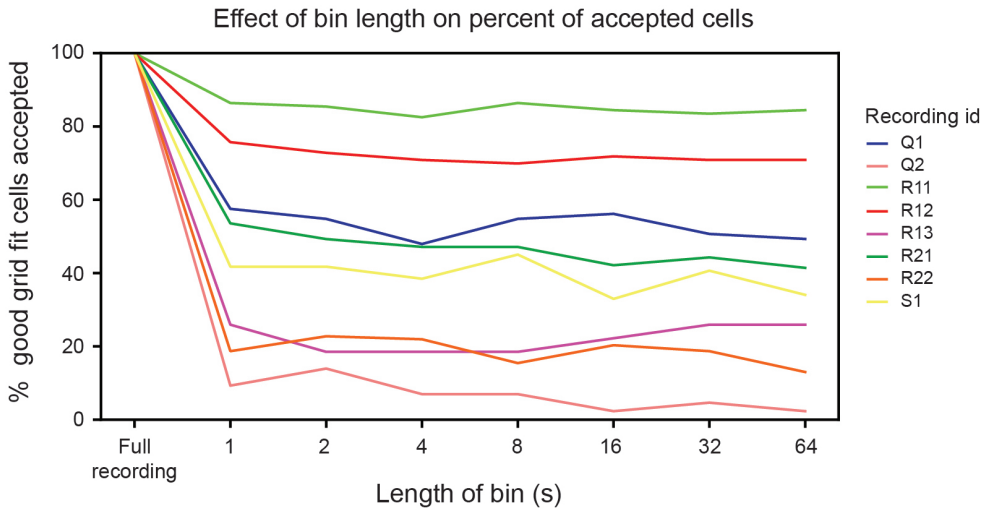

Figure S2: **Bin length does not affect the percent of cells – with good grid fits – that are accepted for analysis.** The percent of cells with good grid fits (i.e., those cells that do not get rejected for having "poor grid fits") that are accepted by not being deemed unreliable, as a function of the size of the bins used in the shuffle analysis. All modules are plotted (each line is colored based on its recording identifier).